Modeling Context Bias

S. P. Velsko, W. C. Thompson

October 1, 2013

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

Modeling context bias

Steve Velsko
Lawrence Livermore National Laboratory

And

William C. Thompson
University of California at Irvine

September 30, 2013

Summary

This report describes a simple model for the effect of domain-irrelevant contextual information on the probative value of expert witness testimony. The model confirms prior observations by Thompson that such information always decreases the probative value. In addition, the effect holds regardless of whether the individual evidence items are conditionally independent or not.

1. Introduction

Context bias refers to the effect of domain-irrelevant information on the testimony of an expert witness.[1] An example of context bias is where a fingerprint examiner declares a match between a suspect's reference fingerprint and a very blurred latent print found at a crime scene *after* the examiner has been informed that a witness has identified the suspect at the scene of the crime, but would have declared the examination "inconclusive" if he did not know about the witness. General discussions of the context bias issue have pointed out that professional forensic examiners often resist efforts to "blind" them to such contextual information.[2] In some cases they argue that contextual information may help them in their examinations, and in other cases they strongly deny that such information affects their judgments. Advocates for blinding often point to experimental demonstrations that context bias affects the judgment of experts in fields such as medicine, and have begun to study the effect directly within the population of forensic examiners.[3,4]

Recently Thompson has developed a simple model based on a Bayesian network to examine the effect of contextual information on forensic expert decision-making.[5] In this network model, an artificial conditional dependence is created between the evidence items proffered by two expert witnesses. Thompson demonstrates that under reasonable assumptions about the effect of the contextual information on the conditional probabilities of the second expert's evidence, the probative value of the his testimony is always decreased if he is exposed to the first witness's conclusions. Thus, if the link between the testimonies of the two witnesses is ignored, the probative value of the combined evidence is overestimated.

This model, while intuitively appealing, presents some formal difficulties. For example, the origin of the contingency matrices that describe the decision process of the second witness is not transparent. Moreover, it leaves open the question of what happens if there is, in fact, an *actual* conditional dependence between the two evidence items. Finally, the formal structure of Bayes's rule for multiple evidence items[6] can engender the false intuition that each expert *must* consider (and thus must be informed of) the other items of evidence in order to render accurate testimony. The analysis in this report attempts to clarify these issues.

Section 2 sets up the formal problem of combining two evidence items using Bayesian reasoning analogous to the model represented by Thompson's Bayesian network. The origin of the decision matrices is clarified by the introduction of a fictitious "meta-expert" who tests the performance of the two experts objectively and reports those results to the Bayesian "jury." The description of the meta-expert's test procedure makes the operational meanings of "conditional independence" and "context bias" transparent.

Section 3 discusses the phenomenology of bias within the model. Consistent with Thompson's observations, the probative value of evidence that is "contaminated" by

another witness's evidence decreases. Moreover, this effect persists even if there is an objective conditional dependence between evidence items. The generality of the model allows us to consider cases where the second expert is <u>mis</u>-informed about the findings of the first.

The final section offers a summary of the results and conclusions of this study. An important connection exists between modeling bias and understanding how to validate the performance of expert witnesses who rely on subjective judgment.   A general rule is that the potential for bias is eliminated when the conditions of an expert's testimony match the conditions under which his performance is validated.

2. Combining evidence in a Bayesian court

We imagine that the jury is a Bayesian machine for converting evidence expressed as categorical assertions into the posterior probability of some hypothesis H against an alternative hypothesis A.[6,7] For simplicity, we assume that there are two pieces of evidence E_1 and E_2, and each piece of evidence has two states T and F (true and false.) For example expert 1 testifies that, after examining the evidentiary material he has been given, in his opinion "$E_1 = T$" or "$E_1 = F$". The jury takes the opined values of E_1 and E_2 and determines the probability of H using Bayes's equation:

$$P(H|E_1E_2) = \left[1 + \frac{P(E_2|E_1A)P(E_1|A)P(A)}{P(E_2|E_1H)P(E_1|H)P(H)}\right]^{-1} \tag{1}$$

We will consider the experts to be "black boxes" whose decisions are characterized by matrices providing the probability that the expert will decide $E_1 = T$ or F, given that the evidentiary material was collected in a situation where the hypothesis or its alternative was true. In other words, the matrices represent $P(E_1|X)$, and $P(E_2|E_1X)$ where X = H or A. We will call these the "decision matrices" for the experts. We can take them to be empirical quantities, based on years of blind testing with mock evidence.

It is convenient to borrow a notion from Cole and imagine a meta-expert who, after each expert states his opinion, provides the jury with the appropriate decision matrix to use in its calculation.[8] How does the meta-expert determine $P(E_1|H)$ and $P(E_1|A)$? He provides the expert with many exemplars of evidentiary materials that have been generated under conditions of H or A respectively. Then he notes the value of E_1 generated by the expert for each exemplar. In this testing process, the meta-expert knows which exemplars are "H" and which are "A" but the expert does not. In other words, the expert is blinded to H and A when he performs his examination.

Similarly, in order to determine the correct decision matrix for expert 2 the meta-expert has to determine how the expert decides the state of E_2 under four conditions: X = H or A, and $E_1 = T$ or F. It is important to understand that, by the same logic used to determine $P(E_1|H)$ and $P(E_1|A)$, expert 2 is blind to <u>both</u> E_1 and X.

If it is found that the empirical $P(E_2|E_1H)$ or $P(E_2|E_1A)$ depend on the state of E_1 even though the expert does not know if $E_1 = T$ or F, we say that E_2 is *conditionally dependent* on E_1.

In the case that E_1 and E_2 were conditionally *in*dependent, the data would show that $P(E_2|E_1 = T,H) = P(E_2|E_1 = F,H)$, and this is equivalent to saying that $P(E_2|E_1H) = P(E_2|H)$. Of course, sometimes one can tell that E_1 and E_2 are conditionally independent simply by asking the question "might an expert decide the state of E_2 differently if E_1 were T or F, *given that he does not know which it is*?"

3. Modeling expert bias

The implication of the previous section is that the meta-expert provides a decision matrix for expert 2 that is based on how expert 2 performs when blinded to the state of both X and E_1. So if expert 2 had been exposed to knowledge about E_1 before rendering a decision, and if this state of knowledge changed his decision matrix, the jury's calculation would clearly be wrong.

There are two potential remedies for this. One is that the meta-expert could determine a decision matrix for the case that the expert has been told the value of E_1 prior to his decision. Note that, to be thorough, the meta-expert must provide 8 different sets of evidentiary exemplars to be evaluated by expert 2: X = H or A, E_1 = T or F, and $\mathcal{E}_1$ = T or F, where $\mathcal{E}_1$ is the variable describing what expert 2 has been told about the state of E_1. However if E_1 and E_2 have been determined to be conditionally independent, E_1 can be left out and only four cases apply. Assuming that this is the case, the meta-expert now reports to the jury the decision matrix $P(E_2|\mathcal{E}_1X)$ where X = H or A, and $\mathcal{E}_1$ = T or F.

The distinction between E_1 and $\mathcal{E}_1$, while subtle, is crucial from a formal point of view. Since it does not make any sense to say that expert 2 is "blinded" to $\mathcal{E}_1$, it should be regarded as a parameter that controls the test conditions under which the decision matrix elements $P(E_2|H)$ and $P(E_2|A)$ are empirically determined. During these tests, the expert is not told whether the exemplars were constructed under H or A, but tests are conducted under conditions where the expert is provided with the information that the value of $\mathcal{E}_1$ is T or F. In this sense $\mathcal{E}_1$ is like any other parameter that might affect the accuracy of the expert's decisions, such as visual acuity or IQ.

Another potential remedy involves the meta-expert providing the jury with a model for estimating $P(E_2|\mathcal{E}_1X)$, based on the known $P(E_2|X)$. A simple intuitive model of how information about E_1 can affect the decision matrix is motivated by the observation that in the extreme "completely biased" case where the expert simply "matches" E_2 to $\mathcal{E}_1$ and ignores all other information, $P(E_2|E_1X)$ can be replaced by a

function $\mathcal{P}(E_2|\mathcal{E}_1)$ which provides the 1:1 mapping between the value of $\mathcal{E}_1$ and the value of E_2 provided by the expert. Thus, a plausible single parameter interpolation between the "no bias" situation (where the expert ignores $\mathcal{E}_1$) and the "completely biased" case is given by the two equations:

$$P(E_2|\mathcal{E}_1 H) = (1-\beta)P(E_2|H) + \beta\mathcal{P}(E_2|\mathcal{E}_1) \qquad\qquad (3a)$$

$$0 \le \beta \le 1$$

$$P(E_2|\mathcal{E}_1 A) = (1-\beta)P(E_2|A) + \beta\mathcal{P}(E_2|\mathcal{E}_1) \qquad\qquad (3b)$$

The parameter β reflects the degree of bias introduced by $\mathcal{E}_1$. The meta-expert might testify that for evidence exemplars similar to the one examined by expert 2 in the case at hand, empirical studies have indicated a certain value for β.

Note that as β increases, the likelihood ratio associated with expert 2's evidence becomes closer to 1, and his testimony becomes more and more irrelevant. Thus, changes to the decision matrix due to influence of $\mathcal{E}_1$ on expert 2 always *reduce* the posterior probability. If the jury (or meta-expert) are not aware that expert 2 was so influenced, and use the unbiased decision matrix $P(E_2|X)$, the posterior probability for H will be exaggerated and the probability of a false positive finding will be increased.

Our discussion so far has assumed that E_1 and E_2 are conditionally independent. Suppose instead that they are not, and the state of E_1 will influence expert 2's decision about the state of E_2 even when he is not aware of what the state of E_1 actually is. If expert 2 hears $\mathcal{E}_1$ before making his decision, precisely the same considerations hold, although the meta-expert now must determine the decision matrix $P(E_2|E_1,X,\mathcal{E}_1)$. The simple model given by equations (3) could be applied equally well, *mutatis mutandis*, and the same conclusions hold.

Conceptual separation of E_1 and $\mathcal{E}_1$ also allows us to consider cases where expert 1 decides one thing but expert 2 hears another. Suppose, for example, that in such a case the jury hears the testimony of expert 2 first, followed by the testimony of expert 1. Expert 2 does not inform the jury or meta-expert that he has heard $\mathcal{E}_1$, and there is no opportunity to reconsider his evidence in light of the actual testimony of expert 1. If $P(E_2|\mathcal{E}_1)$ is diagonal, indicating a tendency to "match" evidence the model represented by equation (3) obviously predicts that this would not effect the likelihood regardless of the value of β.

4. Summary and conclusions

The expansion of $P(E_1,E_2|H)$ into $P(E_2|E_1 H)P(E_1|H)$ does <u>not</u> mean that expert 2 must be informed of E_1 in order to make his decision about E_2. Instead, it means that the "meta-expert" who determines the conditional probabilities that characterize expert 2's performance must be sure that his test materials contain instances where E_1 is varied – but like H, not revealed to expert 2 during testing. If E_1 and E_2 are conditionally (in)dependent, the unknown state of E_1 will(not) affect the performance of expert 2. This makes it clear that if E_1 and E_2 are conditionally dependent, there is some sort of objective causal connection between the two types of evidence.

If expert 2 is informed of E_1 prior to his decision, the decision matrix determined under the testing conditions just described will not be appropriate, and the posterior probability of H is likely be overestimated if that matrix is used in the computation. To determine the correct posterior a decision matrix determined under test conditions where expert 2 is informed by E_1 should be substituted. On the general intuitive grounds that led to equations (3), this modified matrix will result in a lower likelihood ratio for expert 2's testimony, and a lower posterior probability of H.

These conclusions hold regardless of whether or not E_1 and E_2 are conditionally independent. However, it is appealing to say that expert 2 <u>creates</u> a conditional dependence, the causal process being the influence of $\mathcal{E}_1$ on his own mental state. In this regard it is clear that experts whose decisions are based on subjective criteria are most susceptible to this kind of influence. An especially dangerous situation arises *when the expert acts as his own meta-expert* and attempts to inform the jury of both his conclusion and the likelihood ratio associated with it.

What sort of evidence should be withheld from an expert who uses subjective criteria for reaching a decision? The above analysis makes it clear that when there is an objective means for determining the decision matrix under the condition that the expert has been informed of other evidence, then there is no reason to withhold that evidence from him. In the absence of such means, all other evidence should be withheld, even when there is a conditional dependence between his evidence and others. In general, the conditions of his testimony should match the conditions under which his decision matrix (likelihood matrix) was determined.

1. William C. Thompson, "What role should investigative facts play in the evaluation of scientific evidence?", Australian Journal of Forensic Sciences Vol. 43, Nos. 2–3, June–September 2011, 123–134.

2. MJ Saks, DM Risinger R Rosenthal WC Thompson, "Context effects in forensic science: A review and application of the science of science to crime laboratory practice in the United States", Science & Justice Volume 43 No.2 (2003) 77 – 90.

3. Thomas G. Tape, Robert L Panzer, "Echocardiography, Endocarditis, and Clinical Information Bias", Journal of General Internal Medicine Volume 1 **(Sep-Oct).** 1986; pp. 300-304.

4. Itiel Dror, and Robert Rosenthal, "Meta-analytically Quantifying the Reliability and Biasability of Forensic Experts", Journal of Forensic Sciences, July 2008, Vol. 53, No. 4, 900-903.

5. William C. Thompson, private communication, September 2, 2013.

6. Richard O. Lempert, "Modeling Relevance", Michigan Law Review 75(56), 1977, 1021-1057.

7. Franco Taroni, Colin Aitken, Paolo Garbolino, and Alex Biedermann, Bayesian Networks and Probabilistic Inference in Forensic Science", (John Wiley & Sons, West Sussex, England, 2006). Chapter 7: *Aspects of the Combination of Evidence*.

8. Simon A. Cole, "Out of the Daubert Fire and into the Fryeing Pan? Self-Validation, Meta-Expertise and the Admissibility of Latent Print Evidence in *Frye* Jurisdictions", *Minnasota Journal of Law, Science & Technology, Vol. 9 Issue 2; 2008; 453-541.*